

(19)



Europäisches Patentamt

European Patent Office

Office européen des brevets



(11)

EP 0 733 982 A1

(12)

EUROPEAN PATENT APPLICATION

(43) Date of publication:

25.09.1996 Bulletin 1996/39

(51) Int. Cl.⁶: G06F 15/80

(21) Application number: 96104445.0

(22) Date of filing: 20.03.1996

(84) Designated Contracting States:

DE FR GB IT NL SE

(30) Priority: 22.03.1995 IT TO950216

(71) Applicant: CSELT

Centro Studi e Laboratori

Telecomunicazioni S.p.A.

I-10148 Turin (IT)

(72) Inventors:

• Albesano, Dario

Venaria (Torio) (IT)

• Gemello, Roberto

Torino (IT)

• Mana, Franco

Torino (IT)

(74) Representative: Riederer Freiherr von Paar zu

Schönauf, Anton et al

Lederer, Keller & Riederer,

Postfach 26 64

84010 Landshut (DE)

(54) Method of speeding the execution of neural networks for correlated signal processing

(57) The method allows speeding up the execution of a wide class of neural networks for processing input signals evolving slowly through time, such as, for instance, voice, radar, sonar, video signals, and it requires no specialised, costly or hard-to-find hardware. The method requires storing, for the neurons in at least one level of the network, the activation value at a certain

instant and comparing it with the one computed at the subsequent instant. If the activation is equal, the neuron carries out no activity, otherwise it propagates the difference in activation, multiplied by the interconnection weights, to the neurons it is connected to.




Fig. 3

EP 0 733 982 A1

Description

The present invention relates to automatic signal recognition systems and in particular it concerns a method of speeding up the execution of neural networks to process correlated signals.

A neural network is a parallel processing model reproducing, in extremely simplified form, the cerebral cortex organization. A neural network comprises multiple processing units, known as neurons, strongly interconnected by means of links of various intensities, called synapses or interconnection weights.

Neurons in general are arranged along a multi-level structure, with an input level, one or more intermediate levels and an output level. Starting from the input units, which receive the signal to be processed, processing propagates to the subsequent levels of the network up to the output units, which provide the result. A neural network is not programmed, but it is trained by means of a series of examples of the phenomenon to be modelled. Various embodiments of neural networks are described, for instance, in the book by D. Rumelhart "Parallel Distributed Processing", Vol. 1 Foundations, MIT Press, Cambridge, Mass., 1986.

Neural network technology can be applied to many sectors, such as function estimation, audio and video signal processing and classification, automatic controls, forecasting and optimisation, although it still presents some problems stemming from the lack of means that are powerful enough in terms of processing power and speed. It is well known that the execution of a neural network, when it is carried out by emulation on a sequential processor, is very burdensome, especially in cases requiring networks with many thousands of weights. If the need arises to process, in real time, signals continuously varying through time, such as for example voice, video, sonar or radar signals, use of this technology takes on additional difficulties.

The known solutions to this problem have taken three routes.

The first is aimed at decreasing the network size by pruning weights and units, as described for example by R. Reed in "Pruning Algorithms - A Survey" in IEEE Transactions on Neural Networks, Vol. 4, no. 5, 1993. These methods, however, have limited effectiveness since the number of weights and units that can be pruned without incurring a degradation in performance is often very limited.

A second route is based on implementation of the neural network on a VLSI chip, by exploiting its intrinsic potential for parallelisation. This method is, potentially, very promising, but it is not very mature yet. It also entails the use of specialised hardware, which is often very expensive and not easy to integrate with commercial processors.

A third route is the use of specialised hardware of multi-processor type, by distributing the execution of the neural network among various processors. As in the previous case, this possible solution also requires non-

standard hardware, which is costly and difficult to integrate with commercial platforms like personal computers or workstations.

The aforesaid drawbacks are obviated by the method of speeding up the execution of neural network for correlated signal processing, according to the present invention, which allows speeding up the execution of a wide class of neural networks to process sequential input signals evolving slowly through time, such as, for instance, voice, radar, sonar, video signals, and which requires no specialised, costly or hard to find hardware.

In particular, the object of the present invention is to provide a method of speeding up the execution of neural networks for correlated signal processing, as defined in the characterising part of claim 1.

These and other characteristics of the present invention shall be clarified better by the following description of a preferred embodiment of the same, given only by way of non limiting example, and by the enclosed drawings in which:

Figure 1 depicts a Multi-layer Perceptron neural network;

Figure 2 depicts an individual neuron i with its forward connections;

Figure 3 depicts the quantization of the co-domain of the sigmoid transfer function of the neuron.

The idea the method is based upon is the following: since the input signal is sequential and evolves slowly and continuously through time, it is not necessary to compute again all the activation values of all neurons for each input, but rather it is enough to propagate through the network the differences with respect to the previous input. That is, the operation does not consider the absolute neuron activation values at time t , but the differences with respect to activation values at time $t-1$. Therefore at any point of the network, if a neuron has at time t an activation that is sufficiently similar (preferably identical) to that of time $t-1$, then the neuron will not propagate any signal forward. Since the execution of the network is based on the elementary operation of propagating the activation value of each neuron, averaged by the synaptic weights, towards the subsequent neurons, it is readily understood that limiting such activity to just the neurons for which there is an appreciable change in activation level is quite important.

By exploiting this idea, the execution of the neural network can be speeded up by propagating significant activation differences, and this allows saving up to two thirds of execution time in case of speech recognition. This method requires a very small amount of auxiliary memory and it does not entail an appreciable degradation in performance, as it was experimentally verified.

Figure 1 shows a Multi-layer Perceptron neural network like the one described in the already mentioned book by D. Rumelhart "Parallel Distributed Processing". The network input is a signal sampled in time and the

network output are values corresponding to the desired processing, for instance input signal classification.

Each neuron of the intermediate and output levels computes its total input net_i as a summation of the products of the activations o_j of the lower level neurons by the interconnection weights w_{ij} , according to the formula $net_i = \sum_j o_j \cdot w_{ij}$, where w_{ij} is the weight that connects neuron j to neuron i . Then the neuron computes its activation o_i as a non-linear function of the total input according to the formula

$$o_i = 1/(1+e^{-net_i+q_i}),$$

where q_i is a constant value, typical of neuron i . The above non-linear function is a sigmoid, which has an expression of the type $f(x) = 1/(1+e^{-x+\theta})$ and a value between 0 and 1.

Figure 2 shows a single neuron i with its forward connections, along which it propagates the activation differences, and with its memory structures $M1_i$ and $M2_i$ required for the speeding up method. $M1_i$ contains the activation value at time t , $o_i(t)$, as in conventional neural networks, and $M2_i$ the value at the preceding time $t-1$, $o_i(t-1)$. The other neurons in the network also have similar memory structures, for example $M1_k$ and $M2_k$ for neuron k .

Figure 3 depicts the quantization of the set of output values (co-domain) of the sigmoid transfer function of the neuron, with the purpose of quantizing the activation levels of the neurons, thus making it possible to recognise the condition of activation similarity at times t and $t-1$, required in order no signal is propagated. Preferably, the neurons do not propagate any signal when the quantized values at times t and $t-1$ are identical. The elementary operation to speed up execution of the network is thus accomplished.

We shall now describe in detail the method of speeding up the execution of the neural network.

The method requires storing the activation value of a neuron at time $t-1$ and comparing it with the one computed at time t . If the activation values differ by less than a predetermined threshold (preferably, are the same), the neuron carries out no activity, otherwise it propagates the activation difference, multiplied by the interconnection weights, to the neurons it is connected to. Since the activation of a neuron is a continuous non-linear function of its input, usually a sigmoid

$$f(x) = 1/(1+e^{-x+\theta}),$$

it would be necessary to quantize the co-domain of such a function to a pre-determined number of values, in order the condition of activation identity can be actually recognised (Figure 3). Quantization is generally carried out for the neurons of the intermediate level(s), so that the speeding up concerns such level(s).

This quantization of the sigmoid co-domain introduces a sort of quantization noise into the network, but since information is very distributed, this error tends sta-

tistically to be compensated and it does not bring about an appreciable degradation in performance.

Obviously, the number of quantization values must be estimated empirically: the smaller it is, the more the method accelerates; however, it cannot be excessively small to avoid a degradation in performance. In the case of realistic Multi-layer Perceptron networks, with about 50,000 - 100,000 weights, this number can vary from about 25 to about 50.

We shall now describe in detail the execution of the neural network at generic time t , subdivided into the cycles on input neurons j , on intermediate neurons m and on output neurons k .

A) Cycle on input neurons:

1. for all input neurons j the following steps are carried out:

- 1.1 the activation $o_j(t)$ of the input neuron, equal to the j^{th} value of the signal quantized at time t , is defined, after storing the previous activation value $o_j(t-1)$;
- 1.2 for each intermediate neuron m reached from neuron j , the total input net_m is increased by the contribution coming from input neuron j , given by the difference $o_j(t) - o_j(t-1)$ multiplied by interconnection weight w_{mj} .

B) Cycle on intermediate neurons:

2. for all intermediate neurons m , the following steps are performed:

- 2.1 the quantized activation is computed with the formula $o_m^{\text{quant}} = f^{\text{quant}}(net_m)$, where $f^{\text{quant}}(x)$ is the sigmoid quantized on the co-domain, after storing the previous quantized activation $o_m^{\text{quant}}(t-1)$;
- 2.2 if $o_m^{\text{quant}}(t) = o_m^{\text{quant}}(t-1)$, nothing is done for neuron m , and the subsequent neuron is considered, thus accomplishing the basic speeding-up operation; otherwise, the difference is propagated;
- 2.3 for each neuron k reached from neuron m , the total input net_k is increased by the contribution coming from neuron m , which contribution is equal to the difference $o_m^{\text{quant}}(t) - o_m^{\text{quant}}(t-1)$ multiplied by interconnection weight w_{km} .

C) Cycle on the output neurons:

3. for all output neurons k , the exact activation is computed with the formula $o_k = f(net_k)$ where $f(x)$ is the sigmoid.

This method is repeated for all sampled values of the correlated input signal, from start time $t = 1$ to end time $t = T$, where T constitutes the number of elements of the sampled input signal. Before starting to execute the network on a signal, values net_i are initialised to q_i for each neuron in the network, as in the difference propagation mechanism the constant value q of each neuron is defined at the start and it is never altered thereafter.

The main advantages of the method are the following:

- a. optimisation is due to a conceptual modification of the execution of the network and not to a software optimisation, so that the advantage is real and it can be transposed on any hardware platform (even on DSP processor);
- b. the use of additional memory is of the order of the network units, i.e. of the neurons, and not of the order of the weights, so that it is very limited;
- c. by appropriately choosing the quantization step of the co-domain of the sigmoid, it is possible to obtain remarkable reductions in execution time without appreciable losses in performance;
- d. the method can be applied to all cases of sequential input with continuous dynamics and slow evolution (e.g. voice, video, sonar, medical signals);
- e. in the worst case of complete and continuous variation in the input signal, the method converges to the limit case of a standard execution of the neural network, without appreciable additional costs.

It is evident that what has been described is provided purely by way of non-limiting example. Variations and modifications are possible without departing from the scope of the invention.

Claims

1. A method of speeding up the execution of a neural network for correlated signal processing, wherein the network is made of neurons arranged according to a multilevel structure comprising, from a bottom level to a top level, an input level, one or more intermediate levels and an output level, and the neurons in a level are connected to neurons in an immediately upper level by means of links of different intensities (synaptic weights), characterised in that the execution of the neural network at a generic sampling instant t comprises, at least for neurons in the intermediate level or levels, the steps of:

- computing and storing an activation value of the neuron at instant t ;
- comparing the activation value at instant t with the activation value at a preceding sampling instant $t-1$ and computing the difference between said values;

- propagating towards neurons in an immediately upper layer said difference, weighted by the synaptic weights of the links with said neurons in the upper layer, only if said difference exceeds a given threshold.

2. A method according to claim 1, characterised in that said activation values are quantized activation values and in that said comparison is carried out on the quantized values and the activation value difference is propagated whenever the quantized values at instants t and $t-1$ are not identical.

3. A method according to claim 1 or 2, characterised in that the execution of the neural network at instant t comprises the following operations:

a) for all input neurons j :

- a1) defining the activation value $o_j(t)$ of the input neuron, equal to the j^{th} value of the signal quantized at time t , after storing the previous activation value $o_j(t-1)$;
- a2) increasing, for each intermediate level neuron m reached from neuron j , the total input net_m by the contribution coming from the input neuron j , said contribution being equal to the difference $o_j(t) - o_j(t-1)$ multiplied by interconnection weight w_{mj} ;

b) for all intermediate level neurons m :

- b1) computing a quantized activation value with the formula $o_m^{\text{quant}} = f^{\text{quant}}(net_m)$, where $f^{\text{quant}}(x)$ is the sigmoid transfer function of the neuron being considered, quantized on the co-domain, after storing the previous quantized activation value $o_m^{\text{quant}}(t-1)$;
- b2) if the quantized activation value $o_m^{\text{quant}}(t) = o_m^{\text{quant}}(t-1)$, performing no operation for neuron m , and passing to the subsequent neuron;
- b3) increasing, for each neuron k reached from neuron m , the total input net_k by the contribution coming from neuron m , said contribution being equal to the difference $o_m^{\text{quant}}(t) - o_m^{\text{quant}}(t-1)$ multiplied by interconnection weight w_{km} ;

c) for all output neurons k , computing the exact activation value with the formula $o_k = f(net_k)$ where $f(x)$ is the sigmoid.

4. Method according to claim 3, characterised in that it is iterated for all sampled values of the correlated input signal.

5. Method according to claim 3 or 4, characterised in that before starting network execution on a signal, the inputs are initialised by setting $net_i = \theta_i$ for each network neuron, θ_i being a constant value for each neuron, which value is defined at the beginning and is not subsequently changed. 5
6. Method according to any of claims 3 to 5, characterised in that the co-domain of the sigmoid is quantized with a number of quantization steps ranging from 25 to 50. 10

15

20

25

30

35

40

45

50

55

5




Fig. 1




Fig. 2




Fig. 3



European Patent
Office

EUROPEAN SEARCH REPORT

Application Number
EP 96 10 4445

DOCUMENTS CONSIDERED TO BE RELEVANT			
Category	Citation of document with indication, where appropriate, of relevant passages	Relevant to claim	CLASSIFICATION OF THE APPLICATION (Int.Cl.6)
X	PROCEEDINGS OF 1994 IEEE INTERNATIONAL CONFERENCE ON NEURAL NETWORKS (ICNN'94), ORLANDO, FL, USA, 27 JUNE-29 JUNE 1994, ISBN 0-7803-1901-X, 1994, NEW YORK, NY, USA, IEEE, USA, pages 1409-1414 vol.3, XP002007127 GHORBANI A A ET AL: "Training artificial neural networks using variable precision incremental communication" * page 1409, line 1 - page 1410, line 29 *	1	G06F15/80
A	---	2-6	
P,X	IEEE TRANSACTIONS ON NEURAL NETWORKS, vol. 6, no. 6, 1 November 1995, pages 1375-1385, XP000536108 GHORBANI A A ET AL: "INCREMENTAL COMMUNICATION FOR MULTILAYER NEURAL NETWORKS" * page 1375, left-hand column, line 1 - page 1380, left-hand column, line 6; figures 1-4 *	1	
A	PROCEEDINGS OF THE MIDWEST SYMPOSIUM ON CIRCUITS AND SYSTEMS, DETROIT, AUG. 16 - 18, 1993, vol. VOL. 2, no. SYMP. 36, 16 August 1993, INSTITUTE OF ELECTRICAL AND ELECTRONICS ENGINEERS, pages 1085-1088, XP000499764 KWAN H K ET AL: "MULTIPLIERLESS MULTILAYER FEEDFORWARD NEURAL NETWORKS" * abstract; figure 1 *	2,3,6	TECHNICAL FIELDS SEARCHED (Int.Cl.6) G06F

	-/--		
The present search report has been drawn up for all claims			
Place of search THE HAGUE		Date of completion of the search 1 July 1996	Examiner Schenkels, P
CATEGORY OF CITED DOCUMENTS X: particularly relevant if taken alone Y: particularly relevant if combined with another document of the same category A: technological background O: non-written disclosure P: intermediate document		T: theory or principle underlying the invention E: earlier patent document, but published on, or after the filing date D: document cited in the application L: document cited for other reasons A: member of the same patent family, corresponding document	

EPO FORM 1503 (12.92) (P04001)

EP 0 733 982 A1



European Patent
Office

EUROPEAN SEARCH REPORT

Application Number
EP 96 10 4445

DOCUMENTS CONSIDERED TO BE RELEVANT			
Category	Citation of document with indication, where appropriate, of relevant passages	Relevant to claim	CLASSIFICATION OF THE APPLICATION (Int.Cl.6)
A	<p>IEE PROCEEDINGS F. COMMUNICATIONS, RADAR & SIGNAL PROCESSING, vol. 138, no. 1 PART F, 1 January 1991, pages 3-12, XP000179653</p> <p>MURRAY A F: "SILICON IMPLEMENTATIONS OF NEURAL NETWORKS"</p> <p>* page 4, left-hand column, line 34 - right-hand column, line 25; figure 2 *</p> <p>---</p>	2,3,6	
A	<p>PROCEEDINGS OF WORLD CONGRESS ON NEURAL NETWORKS, SAN DIEGO, CA, USA, 5-9 JUNE 1994, ISBN 0-8058-1745-X, 1994, HILLSDALE, NJ, USA, LAWRENCE ERLBAUM ASSOCIATES, USA, pages 11/576-83 vol.2, XP000574437</p> <p>MYUNG WON KIM ET AL: "An asynchronous inter-processor communication based, input recycling parallel architecture for large scale neural network simulation"</p> <p>* the whole document *</p> <p>-----</p>	1	
<p>TECHNICAL FIELDS SEARCHED (Int.Cl.6)</p>			
<p>The present search report has been drawn up for all claims</p>			
Place of search		Date of completion of the search	Examiner
THE HAGUE		1 July 1996	Schenkels, P
<p>CATEGORY OF CITED DOCUMENTS</p> <p>X : particularly relevant if taken alone</p> <p>Y : particularly relevant if combined with another document of the same category</p> <p>A : technological background</p> <p>O : non-written disclosure</p> <p>P : intermediate document</p> <p>T : theory or principle underlying the invention</p> <p>E : earlier patent document, but published on, or after the filing date</p> <p>D : document cited in the application</p> <p>L : document cited for other reasons</p> <p>& : member of the same patent family, corresponding document</p>			

EPF FORM 150 (04/92) (P0401)